CSE 691 COURSE DESCRIPTION FOR SPRING SEMESTER 2025 (3 Credits)
TOPICS IN REINFORCEMENT LEARNING
ASU/SCAI

Prof. Dimitri P. Bertsekas and Dr Yuchao Li

Wednesdays 4:30 ASU time to 7:00 ASU time, BYAC 150

Course Website: https://web.mit.edu/dimitrib/www/RLbook.html

This course will focus on Reinforcement Learning (RL), a currently very active subfield of artificial intelligence. It will discuss selectively a number of algorithmic topics, such as approximation in value and policy space, approximate policy iteration, rollout (a one-time form of policy iteration), model predictive control, large language models, multiagent methods, and their applications to challenging engineering, operations research, and computer science problems.

On the methodological side, our course will be couched on a conceptual framework that centers around two algorithms, which are designed largely independently of each other and operate in synergy through the powerful mechanism of the classical Newton's method. We call these the *off-line training* and the *on-line play* algorithms; the names are borrowed from some of the major successes of RL involving games, such as AlphaZero and TD-Gammon. These algorithms can be implemented in many different ways, and we will emphasize approximate versions of classical Dynamic Programming (DP) algorithms such as value iteration, policy iteration, and rollout.

On the application side, our course will illustrate the RL and approximate DP methodologies within a broad variety of settings involving model predictive and adaptive control, robotics and autonomous systems, large language models, data association, health care, cybersecurity, network infrastructures, and two-person games.

The primary emphasis of the course is to encourage graduate student research in reinforcement learning through directed reading and interactions with the instructors. Prerequisites are a full course on calculus and background in probability.

The course will leverage a series of video lectures, slides, and other material from previous ASU offerings of the course, which are posted at

http://web.mit.edu/dimitrib/www/RLbook.html

Textbooks:

(1) D. Bertsekas, "Reinforcement Learning and Optimal Control," Athena Scientific, 2019.

(2) D. Bertsekas, "Rollout, Policy Iteration, and Distributed Reinforcement Learning," Athena Scientific, 2020.

(3) D. Bertsekas, "Lessons from AlphaZero from Optimal, Model Predictive, and Adaptive Control," Athena Scientific, 2022 (on-line).

(4) D. Bertsekas, "A Course in Reinforcement Learning," 2nd Edition, Athena Scientific, 2024 (on-line). This book will serve as the primary course textbook.

Supplementary material:

(1) Sutton, R., and Barto, A., "Reinforcement Learning," 2nd Edition, MIT Press, Cambridge, MA (on-line). This a valuable resource that approaches the subject from the AI point of view. However, we will not directly use material from this book.

(2) The following survey paper on the relations of reinforcement learning and model predictive control is strongly related to the course: Bertsekas, D., "Model Predictive Control, and Reinforcement Learning: A Unified Framework Based on Dynamic Programming," Published as an IFAC NMPC Preprint, August 2024; slide presentation and video lecture can be found online.

(3) The course's website https://web.mit.edu/dimitrib/www/RLbook.html contains several survey papers and monographs.

Algorithmic Topics:
(1) Introduction to exact and approximate dynamic programming
(2) Approximation in value and policy space
(3) Off-line training, on-line play, and Newton's method
(4) Rollout and approximate policy iteration
(5) Model predictive and adaptive control
(6) Multiagent reinforcement learning
(7) Discrete optimization using rollout
(8) Sequential estimation and Bayesian optimization
(9) Training of feature-based approximation architectures and neural networks

Application Topics:
(1) Robotics and autonomous systems in multiagent environments
(2) Large language models
(3) Inference and optimization of Hidden Markov Models
(4) Data association
(5) Two-person games and computer chess
(6) Infrastructure networks and supply chains
(7) Cybersecurity applications
(8) Health care

Structure:
One 2-hour lecture per week by the instructor, except for the last lecture, which will involve research presentations by student participants. Three-four homeworks (30 percent of the grade), and a research project or term paper (70 percent of the grade).

The first four lectures will introduce the subject and provide a comprehensive overview, helping students focus on a specific research area for their term paper. The subsequent lectures will delve deeper into select topics from the list above, exploring them in greater detail.